(2)

An Investigation of the Impact of Communication Strategies on the Cooperative Behavior of Multiple, Cooperating Autonomous Underwater Vehicles

Phase I Final Report

MSEL Report #90-08

This work is sponsored by the Department of the Navy, office of the Chief of Naval Research under grant no. N00014-89-J-3074. The content of the information does not necessarily reflect the position or policy of the Government.

Prepared for: Office of Naval Research
800 North Quincy Street
Arlington, VA 22217-5000

Prepared by: Roy M. Turner
D. Richard Blidberg
Marine Systems Engineering Laboratory
Marine Programs Building
University of New Hampshire
Durham, NH 03824
(603)862-4600

October 1, 1990

1 Oct. 1990

An Investigation of the Impact of Communication Strategies on the Cooperative Behavior of Multiple Cooperating Autonomous Underwater Vehicles

G-N00014-89-J-3074

Roy M. Turner and D. Richard Blidberg

University of New Hampshire
Marine Systems Engineering Laboratory
Marine Program Bldg.
Durham, NH 03824

90-08

Office of Naval Research
800 North Quincy Street
Arlington, VA 22217-5000

Autonomous Underwater Vehicles (AUVs) hold great promise for use in ocean science and engineering, environmental monitoring, commercial endeavors, and law enforcement and military applications. However, much research remains to be done before AUVs fulfill their promise. Among the many areas which could be targeted for research, we consider the development of multiple cooperating AUV systems to be one of the most important. This paper reports the results of a pilot study aimed at shedding some light on issues related to communication between multiple cooperative AUVs. We were initially concerned with ways to reduce the impact of the inherent uncertainty present in the domain on communication and cooperative problem solving. We began developing a framework for utilizing situation specific schemas representing communication strategies to control communication. While considering the effect of noise and partial knowledge on communication, our attention was drawn to the inadequacy of intention-based communication methods for our domain. Consequently, we began to look at methods of utilizing stereotypical knowledge structures to reduce uncertainty and we have begun to look at ways to interpret a message's content based on the needs of the receiver, not the inferred intentions of the sender.

multiple autonomous underwater vehicles
communication strategies
cooperative behavior,

38

unclassified	unclassified	unclassified	UL

Contents

Summary

Autonomous underwater vehicles (AUVs) hold great promise for use in ocean science and engineering, environmental monitoring, commercial endeavors, and law enforcement and military applications. They can potentially open a window into the vast portion of our planet that remains unexplored beneath the ocean, monitor the quality of seawater over widely-separated areas, repair and resupply submerged structures, and undertake military missions too dangerous for a manned vessel. However, much research remains to be done before AUVs fulfill their promise.

Among the many areas which could be targeted for research, we consider the development of multiple cooperating AUV systems to be one of the most important. Many tasks either cannot be performed by a single agent (AUV) acting alone, or else the task is made much easier (or is accomplished better) when several agents work on it simultaneously. Multiple AUV systems can also greatly increase reliability: if one agent becomes damaged or fails, another can assume its duties. Multiple AUV systems provide an excellent research domain for the emerging field of distributed artificial intelligence (DAI). AUVs operate in an environment fraught with uncertain sensor data, incomplete knowledge of the environment, and unpredictable changes in the world and the AUVs themselves. This means that agents must be able to communicate with one another to share data they discover, to negotiate task allocation, and to coordinate their problem-solving activities. They must be able to flexibly change the way they communicate based on environmental conditions as well as on their own internal states.

This paper reports the results of a pilot study aimed at shedding some light on issues related to communication between multiple cooperative AUVs. We were initially concerned with ways to reduce the impact of the inherent uncertainty present in the domain on communication and cooperative problem solving. While considering the role of communication strategies in reducing the impact of uncertainty and in increasing the effectiveness of communication, our attention was drawn to the fact that a single communication strategy is likely to be inadequate for all of the various kinds of situations an AUV finds itself in. Consequently, we began developing a framework for utilizing situation-specific schemas representing communication strategies to control communication. While considering the effect of noise and partial knowledge on communication, our attention was drawn to the inadequacy of intention-based communication methods for our domain—since it is difficult or even impossible to model another agent's intentions under conditions of partial knowledge. Consequently, we have begun to look at methods of utilizing stereotypical knowledge structures to reduce uncertainty, and we have begun to look at ways to interpret a message's content based on the needs of the receiver, not the inferred intentions of the sender.

Although our efforts in this pilot study have been mostly conceptual, we have begun some programming work to test our ideas. We assembled a research group composed of computer science faculty, graduate and undergraduate students, and staff of the Marine Systems Engineering Laboratory (MSEL) to examine issues related to cooperative distributed problem solving (CDPS), focusing especially on communication and communication strategies. A subset of this group has

begun to design and implement a simulation testbed for CDPS in the AUV domain.

In this paper, we report the results of the pilot study undertaken as Phase I of the project funded by this grant. After discussing our results, we look at directions for work to be undertaken in future phases of our project. These include a more thorough development of the methods developed in Phase I as well as experiments to test our approach, both in our simulation testbed and during in-water experiments aboard MSEL's EAVE AUVs.

1 Introduction

Autonomous underwater vehicles (AUVs) hold great promise for use in ocean science and engineering, environmental monitoring, commercial endeavors, and law enforcement and military applications. They can potentially open a window into the vast portion of our planet that remains unexplored beneath the ocean, monitor the quality of seawater over widely-separated areas, repair and resupply submerged structures, and undertake military missions too dangerous for a manned vessel. They do not suffer from the same drawbacks as human divers or manned submersibles, since they are expendable, can be (potentially) deployed for long periods of time, and are inexpensive, both in terms of initial cost and in terms of cost of operation. Neither do they suffer from the drawbacks inherent in remotely-operated vehicles (ROVs), since they do not require the constant presence of a surface ship, they are not limited by tether length, and they are not susceptible to tether entanglement.

Much research remains to be done before AUVs fulfill their promise, however. Better sensors must be developed, as well as better techniques for interpreting the data from such sensors. Better effectors are needed before AUVs can accomplish tasks requiring manipulation of equipment or the environment. More capable vehicle control software is needed; AUVs will have to be intelligent and capable of adapting their behavior in response to the complex and changing demands of the undersea environment. The use of multiple cooperating AUVs promise an even greater impact in providing solutions to long-standing problems in ocean science and engineering. Research is needed to explore the ways several or many AUVs can cooperate to solve tasks.

Among the many areas which could be targeted for research [see, e.g., Blidberg *et al.*, in press], we consider the development of multiple cooperating AUV systems to be one of the most important. There are several reasons for this. Many tasks either cannot be performed by a single agent (AUV) acting alone, or else the task is made much easier (or is accomplished better) when several agents work on it simultaneously. Examples include:

- A long section of pipe is too unwieldy for a single AUV to pick up and carry by itself; two AUVs, one near each end, could carry the pipe more easily.
- An environmental monitoring task may involve an area too large for a single AUV to cover, either due to the AUV's endurance limitations or because simultaneous measurements are desired from across the area as a whole; several AUVs can cover the area better by working as a team of "underwater satellites" [Blidberg *et al.*, in press].
- Underwater photography, due to the reflective properties of seawater and suspended particles, requires a long baseline between camera and light for good pictures—longer, often, than is feasible on a single vehicle. Two AUVs, one with a camera and one with a light, can cooperate to provide whatever baseline is needed.

Multiple AUV systems can also greatly increase reliability. If one agent becomes damaged or fails, another can potentially assume its duties, allowing the composite problem-solving system to continue its work.

Multiple AUV systems also provide an excellent research domain for the emerging field of distributed artificial intelligence (DAI) [Decker, 1987]. AUVs operate in an environment fraught

with uncertain sensor data, incomplete knowledge of the environment, and unpredictable changes in the world and the AUVs themselves. This means that agents must be able to communicate with one another to share data they discover, to negotiate task allocation, and to coordinate their problem-solving activities. They must be able to flexibly change the way they communicate based on environmental conditions as well as on their own internal states. They communicate over noisy, low-bandwidth channels; AUVs operating in the ocean, for example, must communicate via acoustic telemetry links unless they are very close to one another. Managing the tradeoff between communication bandwidth and processing time is a challenging task.

Before fully-capable multiple AUV systems become a reality, we must address issues in problem solving by individual agents, organization of agents into a composite problem solver (a "cooperative distributed problem solver", CDPS [Durfee *et al.*, 1989]), and communication between agents in such a system.

This paper reports the results of a pilot study aimed at shedding some light on issues related to communication between multiple cooperative AUVs. In particular, we were initially concerned with ways to reduce the impact of the inherent uncertainty present in the domain on communication and cooperative problem solving. The initial goals of our research were:

1. To investigate three aspects of communication between semi-autonomous cooperating systems:
   (a) the knowledge affecting communication;
   (b) the different functional types of communication; and
   (c) the parameters which control the effectiveness of communication.
2. To understand the impact of uncertainty on the cooperative behavior of multiple intelligent systems.
3. To understand how different communication strategies function in the presence of uncertainty and to investigate methods to improve overall performance.
4. To develop methods which allow the performance of various strategies of communication to be assessed.

Our approach to these goals was primarily conceptual. Along the way, our efforts were somewhat re-directed by our conceptual studies. While considering the role of communication strategies in increasing the effectiveness of communication, our attention was drawn to the fact that, as discussed below, a single communication strategy is likely to be inadequate for all of the various kinds of situations an AUV finds itself in. Consequently, we began developing a framework, presented in Section 4, for utilizing situation-specific schemas representing communication strategies to control communication. While considering the effect of noise and partial knowledge on communication, our attention was drawn to the inadequacy of intention-based communication methods for our domain—since it is difficult or even impossible to model another agent's intentions under conditions of partial knowledge. Consequently, we have begun to look at methods of utilizing stereotypical knowledge structures—conventional discourse structures—to reduce uncertainty, and we have begun to look at

ways to interpret a message's content based on the needs of the receiver, not the inferred intentions of the sender.

Although our efforts in this pilot study have been mostly conceptual, we have begun some programming work to test our ideas. We assembled a research group composed of computer science faculty, graduate and undergraduate students, and staff of the Marine Systems Engineering Laboratory (MSEL) to examine issues related to CDPS, focusing especially on communication and communication strategies; much of the discussion below resulted from discussions within the group or between members of the group. Subsequently, a subset of this group began to design and implement a simulation testbed for CDPS in the AUV domain. The initial problem-solving task is underwater photography.[1]

In this paper, we report the results of the pilot study undertaken as Phase I of the project funded by this grant. Our results are of two types, and are reported in separate sections: an examination of issues related to our initial goals, and the development of methods to increase efficacy of communication between agents in a multi-AUV CDPS system. After discussing our results in Sections 3–5, we look at directions for work to be undertaken in future phases of our project. These include a more thorough development of the methods developed in Phase I as well as experiments to test our approach, both in our simulation testbed and during in-water experiments aboard MSEL's EAVE AUVs [Blidberg & Chappell, 1986].

2 Background and Related Work

Cooperative distributed problem solving (CDPS) [see, e.g., Decker, 1987; Durfee *et al.*, 1989] has recently been recognized as a challenging and interesting area of artificial intelligence (AI) research. Single agents, which may have any degree of "local" intelligence, cooperate to carry out tasks that would be more difficult or even impossible to accomplish alone. The entire group of problem solving agents can be considered to comprise a "composite problem solver" that is assigned (or takes on) tasks, then portions them out to its component agents in a manner that best fits the problem and domain.

For the case in which the agents are AUVs, we can make some general observations that could potentially apply to a wide range of CDPS systems:[2]

- Although the group of AUVs comprising the CDPS system must cooperate effectively, they may have a wide range of functional capabilities, communication abilities, and intelligence.
- The group as a whole has a set of tasks to accomplish; although these tasks may be accomplished by a subset of the vehicles or even by a single AUV, they are not independent—there must exist some means of piecing together partial solutions into a globally-coherent solution to the overall task.
- Each vehicle has limited knowledge; in addition, as time progresses, the knowledge common to all vehicles becomes an increasingly small part of the total knowledge

[1]Additional funding for this effort came from NSF grant BCS-8905888 and a grant from the Hubbard Marine Program.

[2]These were also presented in [Blidberg, 1989].

each vehicle has: that is, their local knowledge diverges. Some mechanism must be in place to assure that vehicles have access to information they need when they need it without wasting time and bandwidth sending each item of information discovered.

The first two of these observations suggest that there must exist some means of *organizing* the agents' activities to apportion tasks as well as to ensure global coherence. All of the observations, but especially the last one, suggest that communication is a necessary part of CDPS systems.

Within the set of possible CDPS systems, there are many possible variations along the dimension of agent organization. For example, control may be centralized in a single agent, or it may be distributed among all agents, either in some hierarchical form or in a more egalitarian manner. Organization can be static or dynamic. For example, in work on the scientific community metaphor of organization [Kornfeld & Hewitt, 1981], agents have assigned roles: proposer, critic, etc.; this determines not only the system's control relationships, but also how and to whom the agents communicate. On the other hand, the work of Steeb *et al.* [e.g., 1980] on distributed air traffic control examined schemes for allowing the controlling agent to be selected based on properties of the situation at the time the control decision has to be made. Similarly, the contract net protocol [Smith, 1980] allows agents to have multiple roles (e.g., manager, subcontractor) and to change the roles they play depending on the task being worked on. A benefit of dynamic organization is that the CDPS can adapt, at least to some extent, to the situations in which it finds itself. Organizational issues are of interest to us in this paper to the extent that they are constrained by and in turn constrain communication.

CDPS systems can also vary along the dimension of communication. At one extreme, there are systems that do not communicate at all because communication is infeasible in the situations in which they work. Examples include: covert military and intelligence missions; space missions in which the distance between the agents or the difference between their relative positions renders communication difficult or impractical to carry out in a timely manner; and missions, such as biological experiments, in which communication (e.g., via acoustic underwater link) may adversely impact what is being studied (e.g., dolphin communication).

Cooperation is not impossible, even without communication. Such approaches usually entail each agent having a detailed model of other agents in the environment and of what those agents' behaviors are likely to be. Each agent can then use predictions based on its models to coordinate its behavior with that of the others without the need for communication. However, these approaches are often undesirable due to the high overhead of building and maintaining adequate models of other agents in the environment. Rather than agent A making predictions about what another agent B will do based on its model of agent B—which may include information about B's model of agent A, B's model of A's model of B, etc.—it is often much easier and more reliable for A to simply ***ask*** B what it intends to do. And there are situations in which cooperation without communication is simply impractical. We argue that CDPS in an undersea domain is one such situation. This is due primarily to the inherent uncertainty and unpredictability of the environment. As we discuss below, a characteristic of underwater problem solving is that sensor data is extremely uncertain and

imprecise; consequently it is very difficult to make predictions about the underwater environment, including the activities and internal states of other agents. Cooperation underwater, therefore, requires communication for agents to augment their models of and expectations about other agents.

Communication can serve several purposes during cooperative problem solving. Agents can communicate to establish and maintain an understanding of each other. This relates to the previous discussion of model-based methods of predicting other agents' behavior: communication is used to build and dynamically update the models. Second, an agent can communicate with another to facilitate developing a plan of action to achieve the overall goals of the agents. In cases where all agents are equal partners in problem solving, this communication may look much like human collaborative problem solving and may involve negotiation about what to do; in less equal relationships (e.g., master–slave), one agent may still contribute information to another (e.g., an obstacle or constraint the "master" may not know about). Third, agents can communicate to resolve conflicts, either during the process of planning or while executing the plan. For example, resource limitations may become apparent during plan execution—two or more agents may depend on a resource, yet it may not be sufficient for all. In cases such as this, the affected agents need to communicate to resolve the problem satisfactorily with respect to the overall problem being solved (i.e., to assure global coherence). Fourth, agents can communicate to monitor the progress of the agreed-to plan. For example, an agent can communicate the successes and failures encountered during its portion of the overall plan to other agents. Fifth, agents need to communicate to clarify or verify information. For example, one agent may have better or different sensors than another; consequently, the agent should communicate information its sensors have seen to other affected agents.

- Establish and maintain mutual understanding.
- Facilitate plan development.
- Resolve conflicts during planning or execution.
- Monitor progress of plan.
- Clarify or verify information.
- Synchronize actions of agents.
- Handle emergencies and equipment failures.

Figure 1: Roles for communication during cooperative problem solving.

There are also other roles for communication during problem solving. One such role is to help synchronize actions being carried out by different agents. For example, suppose two AUVs are carrying a long section of pipe between them and need to drop it at a particular location. Instead of each dropping their end whenever they want (which could send an unwary partner to the bottom), or appealing to some global temporal reference system that may not be similarly understood by each agent, the agents should communicate to synchronize their actions: the equivalent of saying "drop it when I say 'now'...now!" Another role of communication occurs during emergencies or

hardware failures. For example, suppose, as is the case on the EAVE AUVs, that an AUV's model of the world is maintained by a different computer from the one on which the planner runs. If the world model computer suffers a failure which results in it being reset, the affected AUV will need to ask others for the information to allow it to rebuild its world model.

Figure 1 summarizes the roles of communication during cooperative problem solving.

Given that the kinds of agents in which we are interested must communicate to cooperate, they can do so in many possible ways. We view an agent as having one or more *communication strategies* (or policies) that control its communicative behavior. There are many different facets of communicative behavior, including: what is communicated, when communication takes place, how information is communicated (e.g., abstract information versus concrete, complete versus partial plans, etc.), and what happens when there is miscommunication or errors in transmission. We concentrate more on communication strategies in Section 3. As with the CDPS' organization, communication strategies can be either static or dynamically selected. We believe that it makes sense for communication strategies to be dynamically selected based on features of an agent's current situation; this ensures that the agent's communicative behavior is appropriate for the context it is in. We discuss this in more detail in Section 4.

Many approaches to communication are based on the receiver inferring the plan of the sending agent based on the messages received and the receiver's model of other agents' knowledge and internal states [e.g., Perrault & Allen, 1980]. For communication in many kinds of systems, this is adequate. However, under conditions of partial knowledge, it is difficult, perhaps even impossible, to maintain an adequate model of another agent. Consequently, plan-based approaches to communication are of less use in a domain such as ours than in many others.

One alternative to plan-based understanding is to use stereotypical or conventional structures instead of inferring plans. For example, scripts [Schank & Abelson, 1977; Bower *et al.*, 1979] have been used to understand stories by filling in missing information, even though the understander maintains no explicit model of the agent telling the story. Conversation memory organization packets (CMOPs) [Kellermann *et al.*, 1989], more flexible structures capturing the shared structure of discourse among speakers, have been suggested to play a role in human communication and have served as the basis for a discourse control system for an advisory computer program [E. Turner & Cullingford, 1989b; E. Turner & Cullingford, 1989a; E. Turner, 1990]. Conversation MOPs allow the merger of convention and intention; as we discuss below, conversation MOPS play a significant role in our approach to multi-AUV communication.

3 Communication Issues

This section discusses issues identified as important to the problem of effective communication between cooperating agents. These include some of the issues identified above.

3.1 Knowledge Affecting Communication

Several kinds of knowledge are important for a problem solver to have during cooperative distributed problem solving. These include:

Knowledge of subject domain. Any agent must obviously have knowledge about its domain, including knowledge about domain objects (e.g., obstacles, ships, etc.), possible environmental conditions (e.g., currents, temperature, etc.) it may encounter, goals or tasks it can reasonably be expected to accomplish, actions, and ways to compose actions into plans. In addition, each agent needs to know about itself and the other agents it may be working with: capabilities and limitations, communication and system characteristics, and so forth. This kind of knowledge impacts communication because it defines the vocabulary of the agents. Attention must be paid to representing this knowledge in such a way that the representation matches the domain and the tasks to be performed. The representation must also be readily accessible to the agents as they need it during their tasks.

Knowledge about the context of current situation. The context of problem solving is one of the most important factors impacting communication between CDPS agents. Whereas domain knowledge is static (leaving aside learning issues), an agent's context constantly changes. The context affects all facets of communication, including with whom to communicate, what to say, how to say it, why, and when it should be communicated.

Several facets comprise a CDPS agent's context. The agent's *environment* determines, to a large extent, what the agents talk about (e.g., "there is an obstacle at (x, y, z)"). The environment can also impact how the agents communicate, for instance if one agent must move to communicate with another agent. An agent's *goals* and *current problem-solving activities* impact communication by motivating what is said (i.e., providing the "why") as well as determining what is said (communication will reflect the needs of the agent's underlying problem-solving activities) and when. *Global goals*—that is, goals to be achieved by the entire CDPS system—impact communication. There must be some realization that there are indeed joint goals or intentions [Levesque *et al.*, 1990]; this assures that agents will communicate, and it influences what is communicated by allowing the agents to assume that other agents will try to help achieve the overall goals. Global goals provide motivation for communicating, and they influence what is said, how, to whom, and when. *Communication channel constraints* can affect communication by restricting the amount of information that can be communicated or the distances over which communication can take place. For example, if the communication channel is acoustic (as it will likely be for an AUV-based CDPS system), bandwidth is low; this can affect communication by increasing the level of abstraction of the information transmitted as well as by decreasing the desirability of using acknowledgments or of repeating information. *Status of self and other agents* also affects communication, since the status of the agents determines who can do what and, hence, whom should be asked to do what. If an agent's status is such that its communicative ability is impaired, this, too, will affect com-

munication with that agent. If an agent's intelligence is lower than other agents operating in the system (e.g., if its control software is simpler) or if its knowledge is restricted, then communication with it may need to be more detailed and concrete. Part of an agent's status is its location, which impacts its viewpoint. This affects communication by introducing (or compounding) the problem of reference: that is, when one agent refers to an object, how does the receiving agent know which object is being referred to? Finally, the *control organization* of the CDPS system impacts communication by defining (or at least biasing) the kinds of messages (e.g., commands versus requests) sent between agents and the agents that are allowed to communicate with one another.

Knowledge about:	Examples
Subject Domain	• objects which may be encountered • possible environmental conditions • other agents that may be present
Context	• current environmental conditions, objects present • goals, current problem-solving activities of agent • global goals and activities • communication channel constraints • status of self, others • current control organization of the CDPS system
Communication	• message formats • how to send messages • discourse structure

Figure 2: Kinds of knowledge affecting communication.

Knowledge about communication. An agent must know how to communicate in order to function as part of a CDPS system. This includes knowledge of agreed-upon message formats for parsing and message composition, knowledge about how to send messages via the selected communications channel, and possibly knowledge about which other agents it is able to communicate with. In addition, knowledge about how to sequence messages is important. Messages should be organized in such a way as to facilitate understanding, reduce bandwidth requirements, and reduce the impact of uncertainty. One approach to doing this is to use shared discourse knowledge structures, such as those described by Kellermann [1989] and E. Turner and Cullingford [E. Turner & Cullingford, 1989b; E. Turner, 1990]. This allows communication to follow common conventions shared by the conversants, facilitating understanding and decreasing the uncertainty involved in communication in a possibly noisy and uncertain domain. Conventions need to be augmented by

paying attention to the intentions of the speaker, since speaker intention strongly influences the the importance and meaning of messages.

Figure 2 summarizes the kinds of knowledge which affect communication.

3.2 Functional Types of Communication

Communication can play many different roles during cooperative problem solving. Some of the roles and functions that we have identified are:[3]

- *Establishing and maintaining knowledge about other agents with whom an agent will* communicate.
- Plan development.
- Execution monitoring, task reporting, and coordination: this enables the composite problem solving system to behave in a manner that facilitates global coherence [Durfee *et al.*, 1989]; it also allows the agents to work out resource and other conflicts as they carry out their tasks.
- Negotiation [e.g., Sycara, 1987; Davis & Smith, 1982]: this includes negotiation during plan development and task execution as well as negotiation related to eliminating inconsistencies arising in two or more agents' knowledge bases.
- Requesting/verifying information.
- Information sharing: at some times during problem solving, an agent may determine that another agent could profit by its information, even if no request was received.
- Emergency-related communication and other warnings: for real-world agents, it is likely that threatening conditions will arise from time to time in the world; consequently, it is important for the agent detecting such conditions to be able to communicate them to the other agents in the CDPS system.

3.3 Parameters Controlling Effectiveness of Communication

One can think of communication effectiveness being impacted by the value of several "parameters"; some of these have been mentioned above, and another, uncertainty, will be discussed below. Here we concentrate primarily on communication parameters and on parameters that can be altered by the agents involved in communication.

One such parameter is the kind of communication channel used, if several are available. For example, an AUV might have both an acoustic telemetry link and a radio link; the radio link, while it has a higher bandwidth and is less noisy, is restricted to use when the agents are at the surface of the sea, while the acoustic link is useful underwater. Factors affected by the choice of communication channel include bandwidth, probability of errors and uncertainty, and amount of noise. These factors, in turn, affect such decisions as whether to use acknowledgments and what to do when there are errors (e.g., whether to repeat messages or not).

Another parameter is the abstraction level of the transmitted information, which can vary along a spectrum from extremely detailed and concrete (e.g., an entire, fully-elaborated plan) to extremely abstract (e.g., a high-level, general description of a plan). A high abstraction level

[3]Some have already been mentioned in Section 2 and are summarized here for completeness.

conserves bandwidth and, in the case of transmitting intentions or orders, allows greater flexibility; however, highly-abstract information may be difficult or even impossible for some agents to interpret and, since the information density is higher, may be more impacted by errors in transmission than more concrete, detailed information. Even when bandwidth is not constraining, agents may elect to use high abstraction levels to conserve transmission and/or parsing time. The best choice of abstraction level, like much else, is likely to vary depending on the circumstances.

The choice of whom to send messages to is another parameter controllable by an agent. The choice here is basically between broadcasting or addressed messages. In the domain of AUVs communicating by acoustic links, broadcasting cannot really be avoided; however, even here, messages can be tagged with addresses allowing agents to ignore messages not addressed to them. The question of which to use is not clear-cut, however: incidental information obtained by "eavesdropping" may help an agent achieve its goals; on the other hand, the proportion of the agent's reasoning resources needed to process all incoming information, rather than only that addressed to it, may outweigh any possible advantages.

An agent also has control of a parameter affecting the overall character of its dialogues with other agents: whether to use shared discourse structures or to engage in unstructured sequences of messages. Structured discourse offers many advantages, among them the ability to engage in coherent, predictable dialogue. This is very useful in domains in which there is much ambiguity or uncertainty, since an ambiguous message can be "fit into" a shared discourse structure: that is, the discourse structure can be used as a context for interpreting the message.

- Kind of communication channel to use.
- Abstraction level of information transmitted.
- Whom to send messages to.
- Whether or not to use conventions to control discourse.
- Static versus dynamic communication strategy selection.

Figure 3: Some parameters affecting communication which are controllable by the communicating agent(s).

An agent also has control over whether it commits to a single communication strategy for the entire CDPS task, or whether it changes the communication strategies it uses dynamically to fit its changing situation. Many previous CDPS systems have taken the former approach [e.g., Davis & Smith, 1982; Steeb *et al.*, 1980; Kornfeld & Hewitt, 1981]. Based on our work, we advocate the dynamic use of communication strategies: it is highly unlikely during a long-duration and/or complex task that a single communication strategy will be both always well-tailored to a situation and applicable across a wide range of situations the agent may encounter. Consequently, we favor an approach in which agents have many communication strategies available and, based on the features of the current situation, the appropriate ones are brought to bear at appropriate times

during the task. This approach is discussed below.

Figure 3 summarizes some of the parameters we have identified which affect communication.

3.4 Impact of Uncertainty on Communication

A focus of our investigation was to identify the impact of uncertainty on communication effectiveness and how that effect can be minimized. A first step towards this was to identify some of the kinds and causes of uncertainty. In the AUV CDPS domain, uncertainty comes in several varieties. *Spatial uncertainty* arises due, in part, to imprecise sensors and incomplete a priori world models. *Temporal uncertainty* can arise from the same causes, but in AUV CDPS systems arises mainly due to the difficulty in synchronizing multiple agents communicating over a channel whose transmission time lapse is not only long (compared to speed of processing) but also varying. If one agent says "I am at (x, y, z) now", when the other agents think the first was at the point depends on how far they are from (x, y, z) themselves—and this, in turn, involves spatial uncertainty. *Reference uncertainty* arises when two agents refer to what they believe to be the same object; however, due to sensor uncertainty, etc., they may or may not actually be referring to the same object. What we call *semantic uncertainty* arises when, for example, a predicate (or word, phrase, etc.) has context-dependent meanings. For instance, a predicate representing "too close" may mean different things depending on whether the objects in question are two AUVs, an AUV and the bottom, or an AUV and a mine. *Causal uncertainty* occurs when an agent cannot predict with any degree of certainty the outcome of an action or event. This is due to weak domain models and/or to partial knowledge. Causal uncertainty applies not only to the environment, but also to other agents. Since we cannot have complete knowledge of another agent in general, given that the agents will be moving about in a highly uncertain world gathering their own sensor data, we cannot model the internal state of another agent with any degree of precision.

Uncertainty impacts communication effectiveness in many ways, as one would expect from the above discussion. Uncertain knowledge can cause errors in parsing or during problem solving, either of which can lead to repetition of messages. Uncertainty can also side-track a dialogue, equivalent to the way "garden path sentences"[4] can side-track human dialogues. Timeliness of information is adversely affected by temporal uncertainty. The content of messages themselves may be uncertain due to noise in the communication medium; this, too, can severely impact communications. Uncertainty with respect to the understanding of the internal state of another agent has the effect of decreasing the usefulness of intention-based, or plan-based, techniques for communication: partial knowledge of another agent's internal state means that an agent cannot predict the other agent's intentions from the messages the agent sends. Reference uncertainty is a problem because it means that an agent may refer to an object, but another agent may not link the reference to the same object in its world model; this can engender a great deal of confusion

[4] E.g., "The old man's glasses were filled with sherry."

and miscommunication. We discuss some ways to mitigate the impact of uncertainty in Sections 4 and 5.

3.5 Communication Strategies

Agents cannot simply communicate at arbitrary times, using messages of arbitrary content and form. If they do, other agents may not be expecting the message and/or may not be able to understand it—or, possibly worse, may misunderstand it and act on the basis of the misunderstood information. Instead, agents need some agreed-upon, shared policies regarding how to communicate. Examples of communication policies include: each AUV communicates only during a particular time slot; an AUV broadcasts its messages rather than sending selectively to another AUV; and messages will have some particular form, perhaps a header indicating type followed by several message fields (i.e., a frame [Minsky, 1975]).

Communication policies, as they have been defined in the distributed AI (DAI) literature, are essentially static, having been agreed upon in advance either by the agents or by the agents' designers. Unfortunately, given the uncertainty inherent in the domain and the changing nature of an agent's environment as it solves problems, static policies are likely to fall prey to two pitfalls. If a policy is general enough to prescribe reasonable communication practices for all situations, it is likely that it will not be specific enough to prescribe optimal or near-optimal practices for any situation. On the other hand, if a policy is tailored to a particular situation or set of situations, it is likely not to be adequate for all situations encountered.

Consequently, an agent needs a repertoire of communication policies from which it can select one to fit its current situation; as the situation changes, it can switch the policy in use. To differentiate this kind of situation-specific communication policy from the static policies in widespread use in DAI, we call them *communication strategies*.

In our view, a communication strategy is a collection of knowledge governing the overall character of communication; it is a *schema* for how to communicate in a particular kind of situation. This last point is very important: *communication strategies are context-dependent*. No single communication strategy will be both appropriate for every situation and sufficiently detailed to be useful in each situation in which it is used. This observation led us to begin developing a framework for dynamically selecting the most appropriate communication strategy to use at each point during cooperative problem solving, discussed in Section 4.

We have begun the process of identifying the kinds of knowledge that should be contained in a communication strategy. In our approach, a communication strategy contains information about what kinds of information in general to send to other agents, when during problem solving to send it, whom to send it to, and how to send it. We discuss this again in more detail in Section 4.1.

3.5.1 What to send?

Communication strategies should allow an agent to know what information it is appropriate to send

to other agents based on the current problem-solving situation. For example, in some situations, it may be appropriate to inform other agents about environmental objects detected via sensors. However, this utilizes bandwidth and causes the other agents to spend time understanding the transmitted messages. In situations when bandwidth or processing time is at a premium, it may be wiser to transmit only messages that are believed to be important to most of the other agents, such as warnings, results of executing a plan, and so forth. The communication strategy in use for a particular situation should help the agent decide what kinds of information it is appropriate to communicate.

3.5.2 When to send?

A communication strategy should also specify when an agent is allowed to send information to another agent. It may make sense for an agent to send any "interesting" results it obtains. For example, when an AUV detects an object via sonar, communicating that result to other AUVs might allow them to modify their path-planning behavior to automatically avoid that object. In other situations, an agent may only want to communicate when predictions are violated which affect the outcome of the plan. If obstacle avoidance is cheaper, in terms of time, than communication (e.g., the communication channel is very slow or there are many AUVs sharing the channel), then it may make more sense to communicate only such things as failure to carry out a task, AUV hardware or software faults, or adverse resource conflicts. There are other situations in which an agent may want to suppress communication altogether. One can easily imagine covert military scenarios requiring silence, but there are also other situations in which communication should be suppressed, at least temporarily. For example, it may be advantageous for AUVs to avoid polluting the acoustic channel for other uses (e.g., shipping, other unrelated AUV teams, etc.) or to avoid disturbance to or detection by what is being studied by the AUV (e.g., cetaceans).

3.5.3 Whom to send to?

Communication strategies should also specify, at least to some extent, with whom to communicate. Basic communication issues, such as whether to broadcast messages or to use some sort of addressing scheme for point-to-point communication, are low-level decisions that should be decided ahead of time as a communication policy. The situation-specific part of the problem is which other agent to communicate with. This may depend partly on the proximity of various other agents—for example, if an AUV needs another sonar view of an object in front of it, it makes sense for it to request this of an AUV that is ten meters away rather than one that is ten kilometers away. The choice of agent to communicate with may depend on the capabilities of the other agent. A simple example is when one AUV has a light and another has a manipulator arm: if the agent wants something lit, it should communicate with the light; if it wants help lifting something, it should communicate with the manipulator. More complex examples involve decisions based on one agent's

perception of others' knowledge or processing abilities; for example, an agent may select another agent because the other agent has always been able to help in past similar situations. Again, since this sort of decision is situation-specific, communication strategies are needed.

Aspect Affected	Issues
What to send	• available bandwidth • processing time of receiver and sender • predicted information needs of other agents • trade-off between cost of processing a message and cost of doing without the information • organization of CDPS system
When to send	• covert versus non-covert behavior desired • trade-off between cost of processing a message and cost of doing without the information • available bandwidth • organization of CDPS system
Whom to send to	• availability of addressing • agent's location • agent's functionality • agent's status • agent's processing capabilities (i.e., "intelligence") • organization of CDPS system
How to send	• available communication channels • bandwidth • processing ability of other agent • constraints on character of way other agent carries out task

Figure 4: Factors affected by an agent's communication strategy.

3.5.4 How to send?

A final use for communication strategies is to determine how to send a message, where "how" refers here to factors such as how abstract the information transmitted should be. In a CDPS situation in which one agent can command or request another to carry out some task, the agent's communication strategy should specify whether to transmit only a description of the task to be accomplished, or to specify the larger plan of which the task is a part. The former saves communication bandwidth and time, while the latter allows some constraints to be placed on how the task is carried out (i.e.,

the other agent should carry out the task in a manner consistent with the overall plan). Though this may seem like a decision that can be made safely before problem solving begins, there are situations in which one technique may be better than the other. For example, suppose that two AUVs are carrying out the task of taking a picture of a feature on the sea floor; one AUV carries a camera, the other a light. If it does not matter how the light-carrying AUV (hence "light") gets to a position as long as the light is shining in a particular way on the object, the camera-carrying AUV (hence "camera") might simply tell the light to move to a given location and orientation, allowing the light considerable freedom in planning its path. However, suppose that the camera knows that the bottom is covered with silt rather than being rocky; consequently, if the light moves between the camera and the object (or too near either one), the conditions for photography will be degraded. In this circumstance, it might make sense for the camera to tell the light its entire plan, which includes "take picture" as a step after the light gets in position. The light can then use this information to reason about its path so as not to violate preconditions (e.g., "water is clear between camera and object") of later steps in the camera's plan. This choice depends on the situation, including both environmental features and such factors as the amount of knowledge the light has (e.g., if the light cannot reason about causes of turbidity, then explicit path constraints would have to be sent rather than the camera's plan). Consequently, the decision is best motivated by a dynamic communication strategy rather than a static policy.

Figure 4 summarizes the factors affected by an agent's choice of communication strategy.

3.6 Performance Methods and Measures

One of our objectives in undertaking this research was to develop methods to evaluate the performance of various communication strategies. Unfortunately, this was a bit premature; much groundwork needed to be prepared before performance metrics could begin to be develop .i. In addition, as part of our research, we identified the need for a mechanism to permit dynamic communication strategy selection; consequently, much of our effort was directed toward developing this method and toward identifying the knowledge contained in communication strategies. Although it was not a realistic objective for the first phase of research, we feel that it is appropriate and important to target as part of future efforts the development of *evaluation methods*: not only performance metrics, but ways of evaluating our overall approach to controlling communication. As a starting point, our evaluation methods will draw from those suggested by Cohen and Howe [1988].

Evaluation will include experimentation aimed at comparing our communication strategies to others published in the AI literature as well as comparing static selection to dynamic selection of our strategies. Partly in support of such experiments, our research group has begun the development of the simulator mentioned previously.

3.7 Summary and New Insights Gained

Research conducted so far has provided many new insights into the problem of communication

and cooperation during multi-agent problem solving. Many of these insights have to do with communication strategies. We believe, based on our examination of communication strategies, that such strategies can be most fruitfully thought of as schemas; that is, as distinct packets of related knowledge that govern the character of communication. Viewed this way, we can see that a strategy must contain a great deal of knowledge, more than is typically discussed in relation to strategies in the literature on CDPS. We touched on this knowledge above: communication parameters, suggestions for discourse structures to use, etc.

Our work also underscored the need for many communication strategies, rather than a single strategy that is used for all situations. Different situations will dictate different kinds of communicative behavior; consequently, it makes sense to have many strategies, each of which is useful in one or more kinds of situations.[5]

Given that an agent needs multiple strategies, a mechanism is needed for selecting the appropriate strategy for the situation. Such a mechanism would use features of the current situation—states of the world, status of the agent, other agents, goals, etc.—to retrieve a strategy that is useful for situations of that type.

Since we are interested in CDPS systems which operate in the real world, agents need a mechanism for switching their communication strategies as dictated by their changing situation. It is likely that part of this mechanism will rely on the mechanism for selecting strategies; however, another part is also needed to notice when a situation has changed enough to warrant changing the communication strategy.

Communication does not exist in isolation; its purpose in CDPS systems is primarily to facilitate cooperative problem solving by the composite system. Consequently, communication strategies must be tightly integrated with the problem solving being carried out by individual agents as well as the overall problem-solving behavior of the overall system.

We have also realized that ways must be identified to assure that compatible communications strategies are selected or agreed upon by all agents engaged in conversation. These methods must assure that as the character of cooperative problem solving changes, so does the set of communication strategies in use by the agents.

A major insight we have gained is that the knowledge contained in communication strategies cannot be obtained without thought for how that knowledge will be used (i.e., without attention to the mechanism for using the knowledge). Similarly, the mechanisms for using communication strategies cannot be completely thought out in isolation from the knowledge that the strategies will contain. Consequently, we have come to realize that the development of communication strategies and the mechanism for using those strategies must occur in parallel.

Other insights we have gained have to do with message structure and the way they are interpreted by agents.[6] Given the partial knowledge an agent may have about the other agents with which it is cooperating, many of the approaches that have been useful in natural language pro-

[5]This idea ties in nicely to our interest and ongoing research in context-sensitive behavior in general.

[6]This is based on work by Elise Turner, co-director of UNH's CDPS research group.

- Communication strategies are (should be) represented explicitly as schemas.
- An agent needs many situation-specific communication strategies.
- A mechanism is needed to select an appropriate strategy at each point.
- A mechanism is needed for switching communication strategies when appropriate.
- Communication strategies must be tightly integrated with the problem-solving activies of the agent.
- A set of cooperating agents must select a compatible set of communication strategies.
- Communication strategies cannot be developed without consideration of the mechanisms which use them.
- Purely intention-based communication is not feasible for CDPS systems involving AUVs: communication between cooperating AUVs needs to follow conventions, and messages need to be interpreted in terms of their importance for the receiver, not the intentions of the sender.
- Partial/uncertain knowledge implies that resolving references (e.g., to objects, etc.) will be a problem for the system.

Figure 5: Some insights gained during Phase I.

cessing research are not as useful in our domain. An agent cannot, for example, rely on correctly interpreting a sender's intentions, since it may not have a reliable model of that agent. Partial knowledge also precludes a sender from too closely tailoring its messages to a particular agent; this is also exacerbated by the fact that several or many other agents may receive the message. What is needed is a way to base understanding more on the recipient's problem-solving needs and less on the perceived intentions of the sender.

We have also determined that reference is likely to be a problem for CDPS systems such as ours. Partial knowledge and uncertainty mean that an agent cannot always be certain that an object to which it refers in a message is the same object that the recipient will think it is referring to. New methods need to be developed to handle reference under highly uncertain conditions.

Figure 5 summarizes the major insights gained while pursuing the goals of Phase I of this project.

4 A Framework for Flexible, Situation-Appropriate Communication

We are in the process of developing a framework for ensuring flexible, appropriate communicative behavior by agents in a CDPS system. Our approach is based on the belief that agents need an explicit representation of their possible communication strategies. This allows them to examine their known strategies, compare what they know about their applicability to the current state of

the world, and then choose the most appropriate one (or ones) for the current situation. In addition, by making the strategies something the agent can examine and reason about, it will be able to use knowledge from the strategy not as an absolute dictum about how it should behave, but rather as an additional source of knowledge about what constitutes appropriate behavior. This allows the agent to tailor its *a priori* knowledge of how it should communicate in a kind of situation—the communication strategy—with its knowledge of the particular situation it is actually in.

4.1 Communication Schemas

In our approach, agents' communication strategies are represented as knowledge structures called *communication schemas* that contain information useful for controlling an agent's communicative behavior in a particular context, or class of situations. Figure 6 lists some of the communication schemas an AUV might need. The way the agent should communicate in each of these kinds of situations is different: for example, use lasers rather than acoustic links, give detailed rather than abstract information, etc. The corresponding schema is what tells the agent how it should communicate.[7]

- Operating with other agents, all equipped with both laser and acoustic communication channels, in clear water and close proximity.
- Operating with those same agents in murky water or at some distance from one another.
- Operating in noisy environments.
- Cooperating with agents that are identical to itself.
- Cooperating with agents that have less intelligence or that are damaged somehow.
- Covert missions, during which no communication should occur.
- Participating in a master–slave organizational structure.
- Participating in a committee organization structure.

Figure 6: Situations for which an AUV would need communication schemas to represent the appropriate strategy to use.

The kinds of knowledge that a communication schema should contain correspond to those kinds mentioned earlier when discussing issues related to communication strategies: what, in general, to communicate; with whom to converse; when, in general, to communicate; and how to communicate. These kinds of knowledge translate into the parts of a communication schema:

1. communication parameter values—these cover some of "when", "what", and "how";
2. knowledge about potential conversants;

[7]Note that these descriptions of situations are not all mutually exclusive, raising the possibility that more than one communication schema might be appropriate for a given situation. We discuss this possibility later.

3. knowledge about how to conduct conversations, that is, knowledge about discourse; and
4. knowledge about what to do in the event of a communication error or misunderstanding.

Figure 7 shows what a communication schema might look like.

Parameters:	channel: acoustic verbosity: moderate acknowledgments: no abstraction level: moderate ...
Conversants:	Self: description of self's capabilites, etc. AUVs: description of other agents, all of roughly the same ability as self
Discourse:	*;; Format is goal: CMOP* Information exchange: CMOP010 Warning: CMOP007 Photography task goals: CMOP002 ...
Error Handling:	general: repeat message ...

Figure 7: A hypothetical (simplified) communication schema representing a strategy to use when communicating during a cooperative photography task.

4.1.1 Parameters

As we discussed earlier, many parameters can affect an agent's communication with other agents. So far, for our domain of cooperating underwater vehicles, we have identified the following parameters, some of which were discussed in Section 3.3:

Appropriate communication channel. Most agents have more than one communication channel they can use. For example, an AUV might have a tether that it uses when communicating with its support vessel at the start of a mission, a radio to use when it is at the surface, an optical communication system for when it is submerged and very near another AUV in clear water, and an acoustic link for when it is far from another agent or when the water is murky [MSEL, 1990]. Selection of the appropriate channel should be automatic and effortless, and the agent should be able to change the channel used as the situation changes. In our approach, communication schemas can suggest the appropriate communication channel to use for the situation the agent is currently in; as the situation changes, which schema is used will also change, and the agent will automatically use the appropriate communication channel, switching as necessary.

Verbosity. An agent's verbosity, both in terms of number of messages and in terms of the verbosity of individual messages, depends on the the situation in which the agent finds itself. In some situations (e.g., when there is noise but the agent has access to a high-bandwidth communication channel), it is better to be verbose; this helps insure that the effect of losses due to noise will be minimized, and, if there is sufficient redundancy in the messages, that message content can be reconstructed even when there is lossage. On the other hand, there are some situations (e.g., when the bandwidth is very low or when engaged on a covert mission) in which verbosity should be kept to a minimum: the increased demand on the recipient's processing time needed to parse the messages, as well as the increased potential of loss of information, is offset in these situations by the corresponding decrease in demand placed on the communication channel bandwidth. In either case, the communication strategy for the situation, represented by a communication schema, should provide information about the appropriate verbosity level. One way that verbosity can vary is in the abstraction level of the messages transmitted.

Whether to acknowledge receipt of messages. Some communication protocols mandate that agents acknowledge receipt of messages. However, this uses bandwidth, sometimes (e.g., in the case of low noise) unnecessarily. By allowing communication schemas to suggest the appropriateness or inappropriateness of using acknowledgments, an agent can make better use of its available communication channel bandwidth.

Kind of information to communicate. An agent must decide, in general, what kinds of information it is going to communicate to other agents, and what kinds of information it can expect to receive. For example, in some situations it may make sense to communicate about objects perceived in the environment; this might help other agents update their world models to maintain an accurate picture of their surroundings. In other situations, we might want the agents to suppress such communications in the hope of decreasing demands placed on the communication channel and on the other agents' processing resources due to receiving messages that may or may not be helpful. Even when a particular kind of thing has been decided upon as appropriate to communicate about, the agent still must decide what, exactly, to say about that thing. For example, an agent may desire to tell another agent to perform a task; should it communicate goals, partial plans, or complete, detailed plans? The answer depends on the situation; for example, if bandwidth is constrained, the recipient is intelligent, and the sender does not care how the task is carried out, then it may make more sense to send a goal rather than a fully-elaborated plan. Since these kinds of decisions are situation-specific, it would seem reasonable to allow the current communication strategy, represented by the current communication schema, to provide information to help the agent make the decisions.

Abstraction level and vocabulary of message. An agent can, as mentioned previously, vary the abstraction level of the information contained in messages it sends. This decreases demand

on the communication channel bandwidth, but increases the demand on the recipient in terms of processing necessary to understand the message. Whether or not the trade-off is worthwhile depends to a great extent on the situation, including how constrained the channel is relative to the recipient's processing resources, the intelligence of the recipient (i.e., if it is insufficiently intelligent, it may not be able to make inferences necessary to understand the abstracted information), etc.

An agent also needs to insure that the vocabulary from which its messages are composed is compatible with other agents' knowledge and processing abilities: i.e., the agents must all "speak the same language". This may not often be an issue, since most CDPS systems involve agents that all are designed by the same researchers and so can be expected to share a common vocabulary. However, for so-called *open systems* [Hewitt, 1986], where agents can come and go as they please, the constituency of a CDPS system will vary and may not always be made up solely of agents which automatically share the same vocabulary. Consequently, for systems such as these, an agent needs to be able to adjust which of possibly several ways of expressing its information it uses. One reasonable way of doing this is to allow the communication schema selected in response to the situation to contain knowledge about what vocabulary to use. This way, the appropriate vocabulary will be automatically selected based on the situation and, as the situation changes, the vocabulary can change, too.

4.1.2 *Conversants*

An agent may be in the presence of many others; it may need to communicate with one, a subset, or all of the agents. Which agents should be communicated with may change as the problem-solving situation changes.

One kind of information that can be used to decide whom to talk to is a description of other agents with whom communication is desirable or advisable. This may be a very simple description, such as "communicate with all"; in this case, the agent would broadcast its messages. However, it may make sense to address messages to particular individuals or groups of agents, if possible. Which is appropriate depends on the situation, including the capabilities and goals of other agents, and may change as the communication strategy changes.

Another kind of useful information is which of several simultaneous conversations to focus attention on at the present time. An agent may be in communication with several other agents simultaneously: sending information to one about the environment, carrying out a task for another, directing a third to perform a task, etc. Some of these conversations may be more important than others; for example, taking a picture of a target for one agent may be more important than sending information about water temperature to another. The communication schema being used by the agent should provide information about the relative importance of kinds of conversations in the current situation. The agent can use this information to select which conversation to focus on.

4.1.3 Discourse

In addition to such factors as verbosity and abstraction level of messages, the way conversations as a whole are structured is another important part of communication. Attention to this aspect of communication is at the *discourse* level.

Conversations can be structured either by the goals of the participants, by agreed-upon (either explicitly or implicitly) conventions, or both. Structuring conversation based on participants' goals, or intentions, is very flexible, but any predictions about the future course of conversation relies on an agent's ability to predict the goals and future goals of the other participants. In noisy, uncertain environments, agents need to be able to predict, at least to some extent, the kinds of information that other participants will be transmitting, since such predictions allow them to disambiguate messages they receive, detect missing messages, and correct misunderstood messages. However, in the presence of noise, uncertainty, and incomplete knowledge, an agent's ability to predict the goals of other agents is severely limited. Consequently, its ability to predict the course of conversation is similarly limited. On the other hand, if participants in a conversation share conventions about what the structure of discourse should be, then the resulting conversation will follow patterns that are predictable by the participants. Humans, for example, tend to follow conventional discourse patterns [Kellermann *et al.*, 1989]. A disadvantage of this approach is that it tends to be inflexible: conversational goals that arise but do not fit into the conventional discourse structure may be ignored.

A compromise position is to use both convention and intention to control conversation. One such approach is E. Turner's work on the JUDIS system, which uses *conversation memory organization packets* to control cooperative problem-solving dialogues between a human and an advisory system [E. Turner & Cullingford, 1989b; E. Turner & Cullingford, 1989a; E. Turner, 1990]. Conversation MOPs (CMOPs) are hierarchical schemas which embody conventions about the way discourse should be structured. JUDIS is guided by the conventions represented by CMOPs, but it is flexible enough to incorporate the intentions of the conversants as well. This allows conversants to predict the general (and in some cases, the specific) future course of conversation while being free to modify the conversation to suit the needs of cooperative problem solving.

We have been and will continue to be investigating the application of E. Turner's approach to CDPS systems in which the agents are AUVs. Obviously, there are no socially agreed-upon conventions in such a system. However, we are optimistic that we will be able to develop discourse conventions for the AUVs to share. This should facilitate communication in the AUVs' noisy, uncertain environment, since they will be able to predict the general kind of message they will receive next; if the actual message does not fit the predictions, then the agent will be alerted to the possibility of error. As importantly, such an approach can potentially limit the bandwidth needed for communication. Just as in human conversations involving scripts [Schank & Abelson, 1977; Bower *et al.*, 1979], all information about a particular subject may not need to be transmitted between AUVs sharing such conventions. Instead, an AUV can send only enough information to

identify the CMOP that is to be used as well as any variations from the conventional structure and content of that CMOP. The receiver, since it also knows the CMOP, can "fill in the blanks".

We do not foresee taking the approach of embedding CMOPs directly in communication schemas. There are two reasons for this. First, the two are different kinds of schemas: conversation MOPs control discourse structure between two conversants, while communication strategies control all facets of an agent's overall communicative behavior in a given situation. Second, an agent may be engaged in conversations, either simultaneously or sequentially, with a variety of other agents; each conversation may be about different topics and may require a different structure of discourse. Instead of directly embedding CMOPs in communication schemas, communication schemas provide suggestions about which CMOPs are appropriate for the current situation. Given these suggestions, the agent can select the one that is most appropriate for the particular conversant and goal, topic, etc.

A communication schema can alter the way a particular CMOP is used to control conversation, in essence modulating the way the (possibly) general-purpose discourse structure is applied so that the resulting conversation more closely fits the current situation. This is done by modifying the CMOP's *MOP-based activation* [E. Turner, 1990]. MOP-based activation is one of two kinds of activation used by JUDIS. At the risk of oversimplification, we can think of *goal-based activation* as representing, to some extent, the contribution of the agent's intentions to the structure of discourse. MOP-based activation, on the other hand, represents the agent's commitment to following conventional discourse structure. In E. Turner's approach, the two types of activation are combined to determine the actual structure of discourse, thus combining contributions of both intention and convention.

The way a communication schema can modify the application of a CMOP is to modify the CMOP's MOP-based activation. This effectively changes the agent's commitment to following convention. In some situations, lowering such commitment may be desirable; for example, if both agents are intelligent and informational messages about the environment are likely to have a large impact on an agent's problem-solving task, then it would likely be best to allow deviations from conventional structure so that the agents can transmit and make use of informational messages as soon as possible. In other situations, it might be better to strengthen the commitment to convention; for example, when dealing with a less intelligent agent, or when the situation is highly uncertain and noisy, a strong conventional component to discourse helps insure that the other agent will understand the messages it is being sent, since it can make stronger predictions about messages given the conventional structure of discourse.

4.1.4 Error Handling

In any realistic scenario, there will be errors in communication. Messages will be garbled or even lost in transmission and agents will misunderstand each other. In situations such as these,

the agents must know how to recover from the errors.[8] How this is done depends on several things. *Message importance* is one factor. If the message is unimportant, then there is no need to take steps to rectify the error; however, if the message is important, then some error-handling actions may be called for. Another factor is the *available bandwidth*. If bandwidth is high and/or utilization is low, then a reasonable thing to do is to retransmit the offending message; however, if bandwidth is at a premium, then other actions may be more appropriate. A third factor concerns *mission goals*; for example, if the agents are involved in a covert mission, then even if bandwidth is high, repeating the message may be inappropriate. And fourth, the *other agent's capability* is a factor impacting error handling. If the other agent is fairly intelligent, it may be appropriate to rephrase a message, possibly at a high level of abstraction, so that the agent can correct its misunderstanding; however, if the other agent is very simple, the more appropriate thing to do may be to retransmit or even to provide a more detailed version of the original message.

There are several possible ways to handle errors and misunderstandings in communication, some of which have already been alluded to. The simplest, with the exception of ignoring the error, is to simply retransmit the message in question. Another method is to transmit, instead of the same message, a message that rephrases it or adds information to help the other agent correct the error or misunderstanding. Another method depends more heavily on the recipient taking the initiative: when an agent detects an error or realizes that it has probably misunderstood a message, instead of simply informing the sender of the problem, it can request specific information to help it correct the problem. For example, suppose an agent sends the message "target is 10m from the rock", but the recipient knows about *two* rocks. In this case, it makes sense for the recipient to request the information that will allow it to disambiguate the first message, i.e., "Which rock?"

We have not focused a great deal on error-handling in our pilot study. Kinds of information necessary and methods to handle errors in our approach will be worked out more fully in future work.

4.2 Using Communication Schemas

Communication schemas affect all facets of communication between agents. They act as a source of knowledge about what, when, and how to communicate with other agents. As such, they are used by any communication functions employed by the agent's reasoning process.

Of primary importance when using schemas to control communication is assuring that the appropriate schema (or set of schemas) is identified as the one to use. A possible selection mechanism is discussed in the next section. However, selection depends on an adequate assessment of the agent's current situation in terms of mission goals, environmental characteristics (e.g., noise, murky water, etc.), and presence and capabilities of other agents [see, e.g., Noble, 1989]. A property of the real world is that it changes, often unexpectedly. Consequently, it is desirable that as the situation changes, the agent re-examines the selected communication schema to assure that it is still the

[8] Error detection is another problem; we have so far given little attention to it.

most appropriate possible one for the new situation: if so, it can continue using the schema as the current communication schema; else, a new schema needs to be selected.

Once selected, a communication schema's information can be used to set parameters affecting the overall character of communication. For example, when an agent finds itself in the situation of communicating via a low-bandwidth acoustic telemetry link, it will select an appropriate schema to provide the strategy for communication in this situation. This schema may suggest that message verbosity be low and of high abstraction, and that message receipt not be acknowledged. This will entail additional work on the part of message recipients to process the messages, and it risks transmission errors, but it conserves bandwidth. If the agent's situation subsequently changes so that it can use an optical communication channel, it would select a different communication schema to guide it in the new situation. This schema's strategy would likely suggest using acknowledgments to reduce undetected transmission errors and would possibly suggest sending more verbose, less abstract messages, both to reduce processing overhead on the part of recipients as well as to provide additional information to help when error recovery is necessary. Notice that once the appropriate schema has been selected, the agent does not have to reason about how to behave; communication behavior is automatic, yet appropriate.[9]

Communication schemas also come into play when an agent must deal with a new conversant. This may happen initially, when problem solving begins—in this case all other agents are "new" conversants—or at any time during problem solving when the agent discovers the presence of a new agent. The agent can use information from its current communication schema about conversants to decide how to deal with the newcomer. For example, the schema may provide a description of a class of agents that fits the new agent and, along with this description, an estimate of such agents' capabilities. If the capabilities predict that the new agent can satisfy some goal currently under consideration, for example, providing information, then the agent may decide to communicate with the newcomer. In addition, the schema may have information about the communicative capabilities of the class of agents; for example, the agents may have only acoustic communication capabilities, or they may not be able to understand and/or deal with highly-abstract goals or information.

Communication schemas also help an agent decide how to conduct a conversation with another agent. This is done by using information contained in the schema about usually-appropriate conversation MOPs to use in the current situation for particular goals, general topics, etc. The agent can use these suggestions to help it select the best CMOP to use to control conversation. In addition, the information about how to set MOP-based activation of the CMOP can be used to tailor the application of the CMOP to the current situation.

It is likely that in a complex CDPS task, an agent may be involved in several conversations simultaneously. When this happens, the agent must decide—often quickly—whom to talk to first. The current communication schema can often provide information about the worth of particular conversations, in terms of their likely impact on problem solver, that can allow the agent to decide

[9]There are times, of course, when the "usually appropriate" thing to do, as suggested by a communication schema, is the wrong thing to do; handling cases such as this is a topic for future research.

which conversation it should focus on at the current time.

Finally, communication schemas are useful when things go wrong. If an error in communication is detected, or if an agent does not understand another, then the schema can be consulted for suggestions of what the appropriate error-handling procedure is for the current situation.

4.3 Organizing and Finding Communication Schemas

An agent will likely need many schemas to control communication, since it will find itself in many different situations as it carries out its problem-solving tasks. A very real problem is how the agent can find the appropriate schema to fit the situation without sacrificing all of the efficiency gained obtained by using schemas due to the effort required to search for a schema. What is needed is a way to quickly retrieve an appropriate schema for a situation based on features of that situation.

One approach to this is to use a *dynamic memory* [Schank, 1982] of the type implemented in the CYRUS program [Kolodner, 1984]. Such a memory is essentially a set of highly-interconnected discrimination nets that are traversed via features of the situation, acting as a conceptual memory for things stored in it. In this case, the nodes of the discrimination nets would be communication schemas. To a large extent, a CYRUS-like memory can also be viewed as an interconnected set of abstraction, or generalization–specialization, hierarchies. In this case, more general communication schemas (e.g., for situations in which the separation between agents is slight) would act to "index" more specific schemas (e.g., for situations in which the separation is slight, but the water is murky). An agent would traverse such a memory in a top-down fashion, stopping when no further specializations are reachable given the current features of the situation. The result would be the retrieval of the most specific schema fitting the current situation.

A memory such as the one described has been used both in the MEDIC schema-based medical diagnostic system to organize procedural and contextual schemas [R. Turner, 1988; R. Turner, 1989b; R. Turner, 1989c; R. Turner, 1989a] and in the JUDIS conversational controller to organize CMOPs [E. Turner & Cullingford, 1989b; E. Turner & Cullingford, 1989a; E. Turner, 1990]. We anticipate using such a memory to organize communication schemas in our approach to CDPS systems.

4.4 Open Issues

There are many open questions to be addressed in the framework described. Some of these are:

- How can we assure that all agents involved in a conversation use compatible communication strategies?
- What should happen when several communication schemas all fit the current situation? I.e., how can information from several schemas be merged?
- When the situation changes, the schema(s) controlling communication should also be switched; how can this be detected and implemented?
- What should the agent do when information from a schema about what is "usually appropriate" communicative behavior is wrong or does not fit well? I.e., what kind of

"from-scratch" meta-reasoning and knowledge does the agent need to decide how to communicate in these circumstances?

- How can communication strategies be identified and elucidated so that they may be stored as communication schemas for an agent to use? I.e., how can we address the issue of knowledge acquisition?
- Can we develop methods to allow the agents to learn new communication schemas and/or modify existing ones based on their own experience?

In addition, much work remains to be done to apply conversation MOPs to the domain of AUV CDPS systems, both from the standpoint of developing conventions to be shared by participating agents as well as from the standpoint of how CMOPs should be used in conjunction with communication strategies to control cooperative problem-solving conversation.

5 Techniques for Communication in Presence of Partial Knowledge

As we discussed above, communication techniques that are reasonable when good models of conversants are available break down when there is only partial knowledge. Good models cannot be created and maintained in general under conditions of incomplete knowledge and uncertainty. Consequently, new methods are needed to ensure the efficacy of communication in CDPS systems in domains such as multi-AUV problem solving.

One technique we have begun to look at involves using discourse structures—CMOPs—to represent conventions shared by all (or most) of the agents involved in communication. These structures provide predictions about messages that may be received, which can allow an agent to infer missing information in a message or disambiguate messages based on the predictions. It allows uncertain or ambiguous messages to be interpreted in the context of knowledge about the ongoing conversation, In order to interface cleanly with the mechanism for using communication schemas discussed earlier, the approach to using CMOPs described by E. Turner [E. Turner & Cullingford, 1989b; E. Turner & Cullingford, 1989a; E. Turner, 1990] will need to be augmented to allow some control of a CMOP's application (e.g., by changing the MOP-based activation) by the currently-active communication schema.

A second technique to decrease the impact of uncertainty is to augment or replace plan-based methods of understanding messages with techniques that rely on interpreting a message's content with respect to the receiver rather than the sender. This entails an agent interpreting a message in the context of its own information needs and problem-solving activities rather than trying to determine what the sender intended to say. This is also mandated by the fact that an agent's message may be received by several or even many different agents, each of which has different uses for the information. For example, it may make sense for one agent to understand another agent's message about an object as effectively warning of an obstacle, while a third agent may best interpret the message as signifying that its target object is nearby. For situations such as this, it

makes sense for the sending agent to send the most general useful message about the object and let the other agents get what information they need from it. Our research group has begun looking at ways of interpreting messages using information local to the receiver.

A third technique we have just begun to look at involves designing message formats, syntax, and semantics to decrease the likely effects of uncertainty. Much more work is needed on this in the future.

6 Status

At the time of writing, we are in the process of implementing a simulation testbed in which to further develop and test our ideas about how to efficiently control communication in noisy and uncertain domains such as multi-AUV CDPS systems. Our simulator's current architecture is shown in Figure 8. Currently it is being implemented to simulate two EAVE AUVs as they cooperate to carry out underwater photography tasks. Ultimately, however, it is extendable to simulate several or many cooperating EAVEs (or other AUVs or agents).

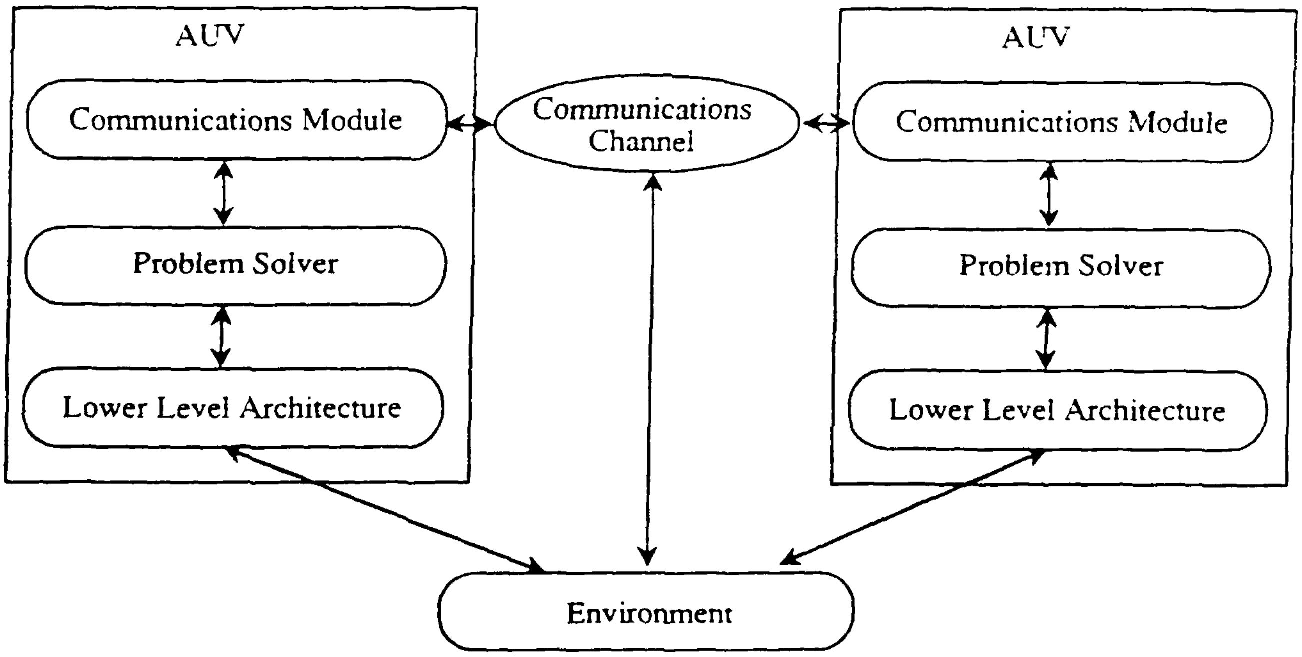

Figure 8: Architecture of the simulator.

The simulator is an object-oriented, graphical system built using Common Lisp and X-windows on Sun workstations. Each AUV is implemented as an object; the environmental model is another object. Communication between the AUVs is through a communication channel object; by changing this object, we will be able to inject noise into messages or model the characteristics of different

communication channels. Within each AUV object, there are currently three modules: low-level architecture, problem solver, and communications. The low-level architecture simulates the bottom three levels of the EAVE software hierarchy (shown in Figure 9). The problem solver and communications modules together correspond to the top layer of EAVE's architecture. The problem solver is currently very simple, essentially just the NOAH planner [Sacerdoti, 1977] coupled with an execution monitor; eventually, the planner will be extended to be an adaptive schema-based reasoner [R. Turner, 1989b; R. Turner, 1989c; R. Turner, 1989a]. The communication module is responsible for much of the kind of communication control discussed above. Communication strategies, however, will ultimately be selected and used by the problem solver and communication modules acting in concert.

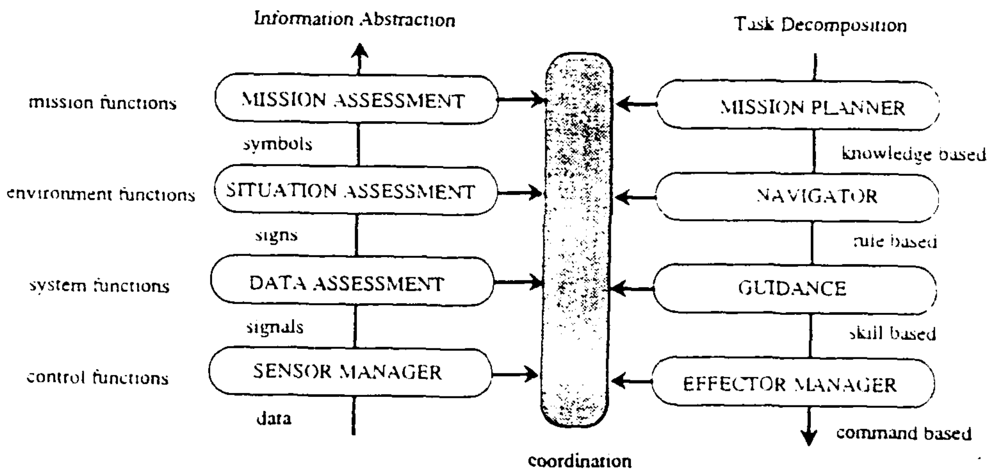

Figure 9: The EAVE software architecture.

7 Future Work

The work done so far has been the first phase of a multi-phase project. In Phase II, we will be concerned primarily with refining and extending the work done in Phase I, as well as evaluating that work by experimentation in a simulation testbed. Phase II goals are:

Develop communication strategies for AUV domain. We have arrived at a first cut at defining the kinds of knowledge that should be contained in communication strategies. This definition, however, still needs refinement. In addition, we need to make our ideas more concrete

and testable by identifying and fleshing out several communication strategies. Consequently, one goal for Phase II is to develop and represent various communication strategies for AUVs to use when involved in a CDPS task. This will allow us to refine our ideas in the context of a real-world domain, one for which expert knowledge is available at our laboratory. As we noted above, the process of developing communication strategies will be undertaken in parallel with the process of refining our model of how to use those strategies.

Refine model of communication strategy use for the AUV domain. During Phase I, we developed a preliminary model of how communication strategies should be brought to bear in a situation-specific manner to control communication between CDPS agents. In Phase II, we will refine this model and detail, by means of a computer implementation, how to apply it in the domain of multi-AUV cooperative problem solving. We will look especially at the issues exposed by our Phase I effort, such as:

- selection of appropriate communication strategies;
- switching the strategies in use as the situation changes;
- ensuring compatibility between strategies used by different agents engaged in a dialogue; and
- integrating the use of communication strategies with the agent's other problem-solving activities.

As a means of approaching the latter issue, we will look at how communication strategies relate to our other research in the area of context-sensitive reasoning.

Identify metrics for evaluation. As mentioned above, Phase I was too early for a serious attempt at evaluation of the approach. However, an important part of Phase II will be to evaluate our approach for controlling communication during cooperative problem solving. One part of the process will be to identify metrics for evaluation. These will need to be applied to compare our approach—both the strategies we develop and our framework for applying those strategies dynamically—to existing approaches, and to compare the static use of our strategies to their dynamic use. Some possible metrics to consider are listed in Figure 10; the ones related to global coherence are derived from the work of Durfee, Lesser, and Corkill [1987a; 1987b].

Evaluate our work. The last goal to be considered here is evaluation of our work; however, we agree with Cohen and Howe [1988] that evaluation should not be left until last, but rather should be an integral part of the entire research process. The kind of evaluation that *will* be undertaken toward the end of Phase II will be experiments using a computer simulation. These experiments will be targeted to:

1. Evaluate our strategies against other strategies/policies in the literature.
2. Evaluate the static use of our strategies against our approach of dynamic strategy selection.

Appropriateness:	• Does our approach increase the appropriateness of an agent's communicative behavior in situations it encounters during problem solving?
Noise and uncertainty:	• Are the effects of noise and uncertainty mitigated? • By how much?
Global coherence:	• Is the information exchanged on average more relevant? • Is the information exchanged more timely? • Is completeness increased? • How can we quantify the improvement of each facet of coherence?
Performance:	• Is the amount of useful effort by each agent increased? • By the system as a whole? (I.e., is the overhead of using dynamic communication strategies worthwhile?) • How does performance degrade under conditions of increased noise, time pressure, low bandwidth, etc.?
Knowledge acquisition:	• How difficult will it be to obtain communication strategies? • How difficult to modify? • Is there a possibility that they can be learned and/or tailored by the program itself?
Resource usage:	• Do the communication strategies promote the effective utilization of resources (e.g., bandwidth, time, etc.)?

Figure 10: Some metrics for evaluating our work.

As part of achieving this goal, the simulation testbed that is already under development will be completed. The simulator will also help us as we go through the process of refining our model.

8 Conclusion

Cooperative distributed problem solving, especially when the agents involved are AUVs, is a challenging and interesting research area with potentially large scientific, engineering, and practical benefits. Of the myriad issues to be addressed in attacking this problem, we have focused our research on those related to communication. Since we are dealing with the underwater domain, this means examining communication under conditions of uncertainty, noise, and partial knowledge of the domain and of other agents.

We discussed many of the issues related to communication in this paper, providing insights we have gained though our pilot study, the first phase of a multi-phase project in the UNH CDPS research group. We have also developed several techniques aimed at mitigating the effects of uncertainty on communication and cooperation during multi-AUV problem solving. One technique involves using explicitly-represented schemas to contain knowledge about communication strate-

gies. These schemas can be retrieved based on the features of the current situation and used to control an agent's communicative behavior in a situation-specific manner. Another technique involves augmenting or replacing plan-based methods of conversation control with methods with a greater emphasis on conversational conventions. Knowledge structures called conversation MOPs are proposed as a way of integrating intention and convention to alleviate some of the problems caused by uncertainty and partial knowledge of other conversants. Another technique involves interpreting messages in the context of what is important to the receiver, not the sender. This again moves away from an intention-based approach to understanding toward a method of understanding more amenable to the constraints of the multi-AUV domain.

We believe that we have made a good start towards elucidating the problems underlying communication during CDPS in the presence of uncertainty, noise, and incomplete knowledge. We have also begun to lay the foundation for an approach for communicating under such conditions that: automatically adapts communicative behavior to fit the current situation; makes use of conventional discourse knowledge to mitigate the effects of uncertainty and partial knowledge; and interprets information in the context of the receiver, not the sender. In future work, we intend to pursue these ideas and to test them, both in simulation and during in-water tests aboard real AUVs engaged in cooperative problem solving.

Acknowledgments

The authors would like to thank the other members of the UNH CDPS research group, especially Elise Turner (co-director), some of whose work is described in this report, Cathy Doman, Peg Eaton, Terri Fisher, Paul Hubbe, and Charu Saxena. Thanks also to Scott McCue for his helpful information on acoustic links and communication issues.

References

Blidberg, D. R. (1989). The impact of communication strategy on the cooperative behavior of multiple autonomous underwater vehicles (AUVs). In *Proceedings of the Sixth International Symposium on Unmanned Untethered Submersible Technology*, pages 331–348, Durham, NH 03824. UNH Marine Systems Engineering Laboratory.

Blidberg, D. R. & Chappell, S. G. (1986). Guidance and control architecture for the EAVE vehicle. *IEEE Journal of Oceanic Engineering*, OE–11(4):449–461.

Blidberg, D. R., Turner, R. M., & Chappell, S. G. (in press). Autonomous underwater vehicles: Current activities and research opportunities. *IEEE Journal of Robotics and Autonomous Systems*.

Bower, G. H., Black, J. B., & Turner, T. J. (1979). Scripts in memory for text. *Cognitive Psychology*, 11:177–220.

Cohen, P. R. & Howe, A. H. (1988). How evaluation guides AI research. *AI Magazine*, 9(4):35–43.

Davis, R. & Smith, R. (1982). Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*.

Decker, K. S. (1987). Distributed problem-solving techniques: A survey. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-17(5):729–740.

Durfee, E. H., Lesser, V. R., & Corkill, D. (1987a). Cooperation through communication in a distributed probelm solving network. In Huhns, M. N., editor, *Distributed Artificial Intelligence*, pages 29–58. Pitman, New York.

Durfee, E. H., Lesser, V. R., & Corkill, D. D. (1987b). Coherent cooperation among communicating problem solvers. *IEEE Transactions on Computers*, C(11):1275–1291.

Durfee, E. H., Lesser, V. R., & Corkill, D. D. (1989). Cooperative distributed problem solving. In Barr, A., Cohen, P. R., & Feigenbaum, E. A., editors, *The Handbook of Artificial Intelligence*, volume IV, chapter XVII, pages 83–147. Addison-Wesley Publishing Company, Inc., Reading, MA.

Hewitt, C. (1986). Offices are open systems. *Communications of the ACM*, 4(3):271–287.

Kellermann, K., Broetzmann, S., Lim, T.-S., & Kitao, K. (1989). The Conversation MOP: Scenes in the stream of discourse. *Discourse Processes*, 12(1):27–62.

Kolodner, J. L. (1984). *Retrieval and Organizational Strategies in Conceptual Memory*. Lawrence Erlbaum Associates, Publishers, Hillsdale, New Jersey.

Kornfeld, W. & Hewitt, C. (1981). The scientific community metaphor. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1).

Levesque, H. J., Cohen, P. R., & Nunes, J. H. T. (1990). On acting together. In *Proceeings of the Eighth National Conference on Artificial Intelligence*, pages 94–99, Menlo Park, CA. AAAI Press/The MIT Press.

Minsky, M. L. (1975). A framework for representing knowledge. In Winston, P. H., editor, *The Psychology of Computer Vision*, pages 211–277. McGraw-Hill, New York.

MSEL (1990). Systems/subsystems investigation for a multi-sensor autonomous underwater vehicle search system. Technical Report MSEL Technical Report #90-01, University of New Hampshire Marine Systems Engineering Laboratory, Durham, New Hampshire 03824. (Also published as Navy EOD Technical Center Technical Report TR-301.)

Noble, D. F. (1989). Schema-based knowledge elicitation for planning and situation assessment aids. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(3):473–482.

Perrault, C. R. & Allen, J. F. (1980). A plan-based analysis of indirect speech acts. *ACJL*, 6(3-4):167–182.

Sacerdoti, E. D. (1977). *A Structure for Plans and Behavior*. Elsevier, Amsterdam.

Schank, R. C. (1982). *Dynamic Memory*. Cambridge University Press, New York.

Schank, R. C. & Abelson, R. (1977). *Scripts, Plans, Goals and Understanding*. Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Smith, R. (1980). The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12):1104–1113.

Steeb, R., Cammarata, S., Hayes-Roth, F., & Wesson, R. (1980). Distributed intelligence for air fleet control. Technical Report WD-839-ARPA, The Rand Corporation.

Sycara, K. (1987). *Adversarial Reasoning in Conflict Resolution*. PhD thesis, School of Information and Computer Science, Georgia Institute of Technology. (Technical report #GIT-ICS-87/26.)

Turner, E. H. (1990). Integrating intention and convention to organize problem solving dialogues. Technical Report GIT-ICS-90/02, School of Information and Computer Science, Georgia Institute of Technology. Ph.D. thesis.

Turner, E. H. & Cullingford, R. E. (1989a). Making conversation flexible. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, pages 932–939, Detroit, MI.

Turner, E. H. & Cullingford, R. E. (1989b). Using conversation mops in natural language interfaces. *Discourse Processes*, 12(1):63–91.

Turner, R. M. (1988). Opportunistic use of schemata for medical diagnosis. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, Montreal, Canada.

Turner, R. M. (1989a). *A Schema-based Model of Adaptive Problem Solving*. PhD thesis, School of Information and Computer Science, Georgia Institute of Technology. Technical report GIT-ICS-89/42.

Turner, R. M. (1989b). Using schemas for diagnosis. *Computer Methods and Programs in Biomedicine*, 30(2/3):199–208.

Turner, R. M. (1989c). When reactive planning is not enough: Using contextual schemas to react appropriately to environmental change. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, pages 940–947, Detroit, MI.